09/763033
JC03 Rec'd PCT/PTO 15 FEB 2001

11pRts



Description

# Fault reversion in a switching device of a communications system

The invention relates to a method for reversion of a fault in an active peripheral assembly of a switching device in a communications system, in particular in an ATM (Asynchronous Transfer Mode) communications system, in which at least one signaled communications link is switched via the active peripheral assembly, and in which connection data for the communications link are stored in the active peripheral assembly in order to handle the communications link. The invention furthermore relates to a switching device for a communications system, in particular for an ATM communications system, having a central control unit for controlling a number of associated peripheral assemblies via which communications links can be switched.

It is known in switching devices of communications systems for communications links to be switched via peripheral assemblies of a switching device, that is to say to handle the communications links with the aid of hardware components of the peripheral assemblies, in particular to set up, to maintain and to end such links. Furthermore, it is known for a number of peripheral assemblies to be controlled via a central switching-internal computer. The central computer transmits, in particular, signals relating to the setting up and clearing of communications links, to the peripheral assemblies.

From ATM communications technology, it is known for permanent and signaled communications links to be set up and to be maintained. The permanent communications links are normally set up and maintained over lengthy time periods by the operator of a communications

system. Permanent communications links are subject to particularly stringent reliability requirements.

Signaled communications links are produced on request by the operator or by a user of the communications system and, as a rule, are maintained for shorter time periods than the permanent communications links. For example, a signaled communications link is ended when a telephone handset is placed on the rest. A typical feature of a signaled communications link is that it is left to the communications system to decide which chain of a number of possible physical path elements for transmission of communications signals will be set up for the signaled communications link.

It is furthermore known for signaled permanent communications links to be set up and to be maintained over a similarly lengthy time period as permanent communications links. Signaled permanent communications links are subject to similarly stringent reliability requirements as permanent communications links. In contrast to permanent communications links, the operator or a user of the communications system specifies only some of the nodes in the communications system and/or only some of the physical transmission sections, via which the communications link is intended to be set up, when setting up a signaled permanent communications link. Apart from this, the setting up of a signaled permanent communications link is left to the communications system. As in the case of signaled communications links as well, signaled permanent communications links require additional connection data and/or connection data of a different type to be stored, in comparison with the data for permanent communications links, in order to handle the communications links.

In order to allow faults in an active peripheral assembly to be reversed, it is known for redundant peripheral assemblies and/or redundant transmission paths to be provided. If a peripheral assembly or a transmission path fails, the communications link or the communications links is or are routed to the redundant peripheral assembly and/or transmission path. In this case, a distinction is drawn between different types of redundancy. For example, peripheral assemblies may have so-called 1+1 redundancy or 1:N redundancy.

In the case of 1+1 redundancy, a passive peripheral assembly is provided for an individual active peripheral assembly and is available exclusively as a standby assembly for that active peripheral assembly. In consequence, the passive, redundant assembly may contain essentially the same hardware and software settings as the active assembly. In practice however, the software settings in particular, that is to say, for example, the connection data stored in the assembly, change at short time intervals. The software for the passive, redundant assembly is thus not reliably at the latest standard, so that it is not possible to switch over from the active assembly to the passive assembly without updating the software and/or the connection data. Furthermore, there is no assurance that the same hardware settings would exist at the same time in both assemblies. For this reason, loss of all the signaled communications links must be expected if it is also possible for some of these signaled communications links to be maintained. Furthermore, when switching over from an active transmission path to a 1+1 redundant transmission path, it may be necessary to switch a number of peripheral assemblies which each have 1+1 redundancy. This increases the probability that the majority, or all the signaled communications links, will be lost.

In the case of 1:N redundancy, there is only one joint redundant assembly where the number of assemblies is N > 1. Presets in the redundant assembly, which speed up switching from one of the N active assemblies to the redundant assembly, which ensure that existing signaled communications links are maintained, thus cannot be carried out, or can be carried out only to a limited extent, according to the known method.

When software faults occur in an active peripheral assembly, it is known for the entire assembly to be reset and to be restarted. In the process, all the connection data for signaled communications links are deleted, so that all these communications links are terminated.

The present invention is based on the object of specifying a method for reversion of a fault in an active peripheral assembly of the type mentioned initially, in which stable signaled communications links can be reliably maintained after the occurrence of the fault. A further object of the invention is to specify a corresponding switching device.

The objects are respectively achieved by a method having the features of claim 1, and by an apparatus having the features of claim 10. Developments are the subject matter of the dependent claims.

With regard to the method, the connection data for the at least one communications link, which data are stored in the active peripheral assembly, are also stored elsewhere in a redundant manner. After the occurrence of the fault, the connection data which are stored elsewhere in a redundant manner are used in order to continue to handle the communications link. The storage of two up-to-date sets of connection data ensures that the connection data are still available, at least once, after the failure or after the occurrence of the fault in the active peripheral assembly. Provided the redundant set of connection data, or at least one of the redundant sets of

connection data, is or are undamaged after the occurrence of the fault, the undamaged data set can be used to continue to handle the communications link. It is thus possible to maintain the at least one signaled communications link, provided it is still stable after the occurrence of the fault. In order to preclude the possibility of the set of data which is stored in the faulty peripheral assembly being damaged and thus endangering the maintenance of the at least one communications link when it is used further, the connection data stored in the faulty peripheral assembly are preferably deleted.

The method according to the invention has the advantage that signaled communications links can be maintained with the same reliability as permanent communications links, which may be unstable in the same way as the signaled communications links after the occurrence of the fault, or can no longer be maintained owing to the fault.

The term peripheral assembly means an assembly or unit of a switching device which is directly involved in the setting up and/or in the maintenance of a communications link. The term peripheral assemblies covers, in particular, interface assemblies (Line Interface Cards LIC), which form an interface to transmission lines outside the switching system, multiplexer units (for example Statistical Multiplexing Units SMU) which connect a number of interface assemblies to a switching network, switching networks (ATM Switching Networks ASN) which produce the correct output from a number of possible outputs in response to signals arriving at them. In particular, the assemblies have a large number of elements which carry out functions on the assembly, but at least a memory area for storing the

connection data and an element which is directly involved in the communications link. The invention can also be used in other communications system, for example in STM (Synchronous Transfer Mode) communications systems.

In a communications system, the term switching device means a device which allows various switching-system-external and/or switching-system-internal transmission paths or transmission lines to be connected to one another or to be enabled for the purposes of setting up a communications link.

The term communications link means a link for connection of any desired type via which signals can be transmitted within the communications system or beyond the boundaries of the communications system.

The connection data are preferably stored in a memory device which is central for the number of peripheral assemblies, before the fault occurs. In particular, the memory device is part of a central switching computer of the switching device. In this case, the switching computer may, for example, update that copy of the connection data which is stored in its memory device whenever it informs the active peripheral assembly of changes which relate to the connection data.

However, on the other hand, as well, the central memory for the connection data for a number of peripheral assemblies has the advantage that one common memory unit or memory device is sufficient, and that the data administration can be carried out in an effective manner.

In one development, a redundant passive peripheral assembly, in which the connection data are stored in a redundant manner, is provided for the active peripheral assembly. In contrast to the prior art, the passive peripheral assembly is used to store a set of connection data which is just as up-to-date as that in

the active peripheral assembly. This firstly creates the possibility simply of switching over to the passive peripheral assembly after the occurrence of the fault in the active peripheral assembly while, on the other hand, it is possible to transmit the connection data from the passive peripheral assembly to the active peripheral assembly, for example if the connection data stored in the active peripheral assembly are faulty, or there is a possibility of such data being faulty.

If a fault occurs in the software of the active peripheral assembly, the connection data which are stored elsewhere in a redundant manner can be transmitted, in a preferred manner, to the peripheral assembly which is still active after the occurrence of the fault. Faults in software covers not only faults in programs, but also faults in stored data to which access can be made while a program is running.

In many fault situations, continuation of operation of the active peripheral assembly represents the fastest and most reliable option for maintaining stable communications links without considerable interruptions. Mostly, hardware settings still exist in the active peripheral assembly in order to maintain a communications link after the occurrence of the fault. This is preferably checked after the transmission of the connection data stored elsewhere in a redundant manner, or at least after transmission of some of these connection data items. During the check it is also, for example, possible to confirm whether the connection data stored elsewhere in a redundant manner have been changed as a consequence of their having been updated in the meantime, so that corresponding corrections or changes are required to the hardware settings. One possible reason for this occurs when a message that the communications link is intended to be ended arrives at a central switching computer in the meantime.

In one preferred refinement of the method, the connection data to be transmitted remain stored at the other location, that is to say they are transmitted in the form of a copy. This also applies in particular to the situation in which, after the occurrence of the fault, the previously active peripheral assembly becomes passive and a redundant assembly is used as the active peripheral assembly, to which the connection data stored elsewhere in a redundant manner are transmitted. This method variant makes it possible to maintain stable communications links in particular when hardware faults occur in the previously active peripheral assembly.

In one particularly preferred development of the method, the transmission of the connection data to the peripheral assembly which is active after the occurrence of the fault is interrupted, or is not started until later, in order to allow the setting up of new communications links. In this case, expediently, the connection data to be transmitted are transmitted in blocks to the active peripheral assembly. One major advantage of this development is that, as soon as it is confirmed which peripheral assembly will be taking over or continuing the active operation after the occurrence of the fault, new communications links can be set up, and in this case it is possible to use the same method as that for fault-free operation.

With regard to the apparatus, the object of the invention as stated above is achieved in that the central control unit of the switching device has a data memory in which connection data can be stored from signaled subscriber connections which are switched via the associated peripheral assemblies, and in that a transmission unit is provided for reading and transmitting the connection data to the associated peripheral assemblies. Advantages and developments result from the above description of the method according to the invention, and its developments.

Exemplary embodiments of the invention will now be described with reference to the attached drawing. However, the invention is not limited to these exemplary embodiments. The single figure of the drawing, which is annotated Figure 1, shows:

a switching device having two 1+1 redundant interface assemblies during reversion of a fault.

The switching device EXC (shown in Figure 1) of a communications system has two interface assemblies LIC1, LIC2 which are 1+1 redundant with respect to one another. A large number of communications links are set up, maintained and/or ended via the respective active interface assembly LIC1, LIC2. For this purpose, connections (which are not shown) of the interface assemblies LIC1, LIC2 are connected to external subsections (which are likewise not shown) of transmission lines for transmission of communications signals. Within the switching system, the interface assemblies LIC1, LIC2 are connected to additional assemblies (which are also not shown) of the switching device EXC. The communications signals arriving at and/or departing from the active interface assembly LIC1, LIC2 can be transmitted on defined sections of communications lines via these assemblies, for example multiplexer assemblies, switching network assemblies and other interface assemblies.

The fault reversion described in the following text can also be carried out in a corresponding manner for the said types and further types of peripheral assemblies for a switching device. In a first exemplary embodiment for fault reversion, a software fault has occurred in the peripheral assembly LIC2. In order to circumvent the fault, the interface assembly LIC2 is reset, and all the data that are not stored in permanent memories are lost.

This effect is desirable since it ensures that faulty data are deleted. The effort for precise localization of faulty components of data which are not faulty overall is generally considerably greater than this.

After the resetting of the interface assembly LIC2, basic data are first of all transferred or transmitted from a database manager DBMS of a central switching computer MP to a database CDB of the interface assembly LIC2. In the process, it is possible to transmit, in particular, data for programs which are used to carry out the functions of the peripheral assembly LIC2. Alternatively or additionally, a permanent memory, for example an FEPROM, is provided in the peripheral assembly LIC2, in which program data are stored and from which, after resetting, data are copied to a volatile memory, for example a RAM, which contains the database CDB. It has already been proposed that, in method step a, data be transmitted from permanent communications links to the interface assembly LIC2 which are stored in a redundant manner in the central switching computer MP. It is also known, once method step a has been carried out, for hardware settings relating to the connection hardware ASIC of the interface assembly LIC2 to be checked for consistency with the data of the permanent communications links, and for any corrections which may be necessary to the hardware settings to be carried out. The connection manager COH of the interface assembly LIC2 is involved in this activity, in that it reads the data from the database CDB and carries out the appropriate test steps.

Now, in method step b, connection data, which are stored in a redundant manner in the central switching computer MP, for the signaled communications links are transmitted to the interface assembly LIC2, which is being or has been maintained with the aid of the connection hardware ASIC in the interface assembly LIC2.

In the process, it is possible for individual signaled communications links, or for a number of signaled communications links, to be interrupted or terminated as a result of the fault which has occurred. In method step b, a switching data manager RHS in the central switching computer MP transmits the connection data for the signaled communications links to the connection manager COH in the interface assembly LIC2. In the process, the connection data are transmitted in data blocks of predetermined length. It takes the connection manager COH a period of 5ms, for example, to receive a data block, with connection data being received for about 30 signaled communications links for the interface assembly LIC2. After reception of the data block, the connection manager COH starts to check the hardware settings of the corresponding communications links, which are carried out in the connection hardware ASIC (method step d). It takes the connection manager COH, for example, about half a second to read the connection data and to check them for consistency with the hardware settings for about 30 communications links.

Method step b' is carried out in parallel with method step b). In method step b', the connection manager COH is informed by the connection data manager RHS when a request to set up a new communications link is present in the central switching computer MP. Before the transmission of the first data block in method step b, or between the transmission of two data blocks, the connection manager COH reacts to the request by setting up a new communications link on the basis of the connection data which it receives from the connection data manager RHS together with the information about the presence of the request. The new communications link may be a signaled communications link or a non-signaled communications link. Corresponding procedures are used in the presence of requests to set up a number of new

communications links. In this case, the corresponding data can be transmitted individually or in blocks to the connection manager COH. The requests for new communications links preferably have priority over the checking of the hardware settings of already existing communications links, so that they are processed with priority by the connection manager COH. Furthermore, however, a time interval of predetermined length is preferably provided after which, at the latest, the connection manager COH again receives connection data from existing communications links, or continues with checking the hardware settings of communications links whose connection data it already has. For example, during operation of communications systems in practice, it is accepted that the setting up of requested new communications links in a peripheral assembly will be delayed by around 500ms. As described above, the connection data for approximately 30 already existing communications links, for example, can be read and the corresponding hardware settings checked within 500ms. The checking of the hardware setting is annotated in Figure 1 by the reference symbol d, and the resetting of communications links in the connection hardware HSIC is annotated by the reference symbol d'.

Before, during and after reception of connection data for already existing communications links and/or reception of connection data of communications links to be set up as new, a start pulse for a time-interval measurement for connection-duration-dependent charges is in case issued, in a development of the invention. If the connection manager COH has the respective connection data for such a communications link, it sends the start pulse to the charge meter TM which is provided in the interface assembly LIC2. The transmission of the charge-meter start pulse for already existing communications links is illustrated as method step c, and the corresponding

transmission of a start pulse for a communications link which is to be set up as new is shown as method step c'.

In yet another development of the exemplary embodiment described with reference to Figure 1 for the method according to the invention, the central switching computer MP transmits a request to the connection manager COH to terminate an existing communications link, in a method step (which is not shown) after method step a. However, in the development of the exemplary embodiment, the charge-meter start pulse according to method step c has not yet been transmitted from the connection manager COH to the charge meter TM at this time. The connection manager COH thus does not confirm the request, and does not act on this request either. This ensures that the charge meter TM can determine the connection-duration-dependent charges correctly and does not receive a charge-meter stop pulse before receiving a charge-meter start pulse for an existing communications link. Since the central switching computer MP have not received any confirmation from the connection manager COH, it repeats the transmission of the request to terminate the communications link, preferably at time intervals which are fixed in advance. In the meantime, for example between the first and the third transmission of the request, the connection manager COH will have received the connection data for the relevant communications link, and will have sent a corresponding charge-meter start pulse to the charge meter TM. This thus acknowledges the third request with a confirmation, sends a charge-meter stop pulse to the charge meter TM, and terminates the communications link by making appropriate hardware settings in the connection hardware ASIC.

In the exemplary embodiment described so far, it has been assumed that there has been a software fault in the interface assembly LIC2. The same method, together with its development, can, however, be applied, for

example, to a situation in which a hardware fault has occurred in the previously active interface assembly LIC1 and the 1+1 redundant interface assembly LIC2 has taken over the active function from the interface assembly LIC1. Furthermore, the same refinements of the method can be used in situations in which there is no redundant component or in which 1:N redundancy exists. Overall, a universal method is thus available for reversion of a fault in an active peripheral assembly in a switching device, which allows all those communications links which are stable after the occurrence of the fault to be maintained while having to accept, at worst, only short interruption times. Furthermore, the universal method ensures the setting up of communications links as new with the least possible time delay while reliably deleting all possibly faulty data in the assembly.